

Claims

What is claimed is:

1. A method of determining the location of an Internet host using a first computer system, comprising:

obtaining route information relating to first and second network paths between a host IP address associated with the Internet host and the first computer system and a second computer system, respectively, wherein the first network path comprises the first computer system, the Internet host, and at least one intermediate network node, the second network path comprises the second computer system, the Internet host, and at least another intermediate network node, and wherein the route information comprises a plurality of router labels associated with the host IP address and one of the at least one intermediate network node and the at least another intermediate network node;

extracting a first location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

consulting a data store comprising at least one data set having location codes and corresponding location information;

obtaining first location information from the data store corresponding to the first location code associated with the one of the Internet host and the intermediate network node proximate the Internet host; and

providing a first location estimate of the location of the Internet host according to the first location information from the data store corresponding to the first location code.

2. The method of claim 1, further comprising:

determining a delay time associated with at least one of the first and second network paths; and

selectively correcting the first location estimate according to the delay time associated with the at least one of the first and second network paths.

3. The method of claim 2, further comprising determining a confidence metric representative of the accuracy of the first location estimate according to the delay time.

4. The method of claim 1, further comprising:
determining a confidence metric representative of the accuracy of the first location estimate; and
selectively providing the first location estimate of the location of the Internet host if the confidence metric exceeds a threshold.

5. The method of claim 1, further comprising:
obtaining second route information relating to the second network path between the host IP address and the second computer system wherein the second route information comprises a plurality of router labels associated with the host IP address and the at least another intermediate network node;
extracting a second location code from the second route information corresponding to a second router label associated with one of the Internet host and the at least another intermediate network node;
obtaining second location information from the data store corresponding to the second location code;
providing a second location estimate of the location of the Internet host according to the second location information from the data store corresponding to the second location code; and
correlating the first and second location estimates to provide an improved location estimate of the location of the Internet host.

6. The method of claim 1, wherein the location code comprises one of a city code, and airport code, and a country code, and wherein obtaining the route information comprises using a traceroute tool.

7. A software tool for determining the location of an Internet host using a first computer system, comprising:

a route trace component adapted to obtain route information relating to first and second network paths between a host IP address associated with the Internet host and the first computer system and a second computer system, respectively, wherein the first network path comprises the first computer system, the Internet host, and at least one intermediate network node, the second network path comprises the second computer system, the Internet host, and at least another intermediate network node, and wherein the route information comprises a plurality of router labels associated with the host IP address and one of the at least one intermediate network node and the at least another intermediate network node;

an extraction component adapted to extract a first location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

a data store comprising at least one data set having location codes and corresponding location information; and

an estimation component adapted to obtain first location information from the data store corresponding to the first location code associated with the one of the Internet host and the intermediate network node proximate the Internet host, and to provide a first location estimate of the location of the Internet host according to the first location information from the data store corresponding to the first location code.

8. A computer-readable medium having computer-executable instructions for:

obtaining route information relating to first and second network paths between a host IP address associated with an Internet host and a first computer system and a second computer system, respectively, wherein the first network path comprises the first computer system, the Internet host, and at least one intermediate network node, the second network path comprises the second computer system, the Internet host, and at least another intermediate network node, and wherein the route information comprises a

plurality of router labels associated with the host IP address and one of the at least one intermediate network node and the at least another intermediate network node;

extracting a first location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

consulting a data store comprising at least one data set having location codes and corresponding location information;

obtaining first location information from the data store corresponding to the first location code associated with the one of the Internet host and the intermediate network node proximate the Internet host; and

providing a first location estimate of the location of the Internet host according to the first location information from the data store corresponding to the first location code.

9. A system for determining the location of an Internet host, comprising:

a first component operating in a first computer system to obtain route information relating to first and second network paths between a host IP address associated with the Internet host and the first computer system and a second computer system, respectively, wherein the first network path comprises the first computer system, the Internet host, and at least one intermediate network node, the second network path comprises the second computer system, the Internet host, and at least another intermediate network node, and wherein the route information comprises a plurality of router labels associated with the host IP address and one of the at least one intermediate network node and the at least another intermediate network node;

a second component operating in the first computer system to extract a first location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

a third component operating in the first computer system to consult a data store comprising at least one data set having location codes and corresponding location information;

a fourth component operating in the first computer system to obtain first location information from the data store corresponding to the first location code associated with the one of the Internet host and the intermediate network node proximate the Internet host; and

a fifth component operating in the first computer system to provide a first location estimate of the location of the Internet host according to the first location information from the data store corresponding to the first location code.

10. Geographical location estimate data associated with an Internet host, the estimate data resulting from a process of:

obtaining route information relating to first and second network paths between a host IP address associated with the Internet host and a first computer system and a second computer system, respectively, wherein the first network path comprises the first computer system, the Internet host, and at least one intermediate network node, the second network path comprises the second computer system, the Internet host, and at least another intermediate network node, and wherein the route information comprises a plurality of router labels associated with the host IP address and one of the at least one intermediate network node and the at least another intermediate network node;

extracting a first location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

consulting a data store comprising at least one data set having location codes and corresponding location information;

obtaining first location information from the data store corresponding to the first location code associated with the one of the Internet host and the intermediate network node proximate the Internet host; and

providing a first location estimate of the location of the Internet host according to the first location information from the data store corresponding to the first location code.

11. A method of determining the location of an Internet host using multiple computer systems, comprising:

obtaining route information relating to a plurality of network paths between a host IP address associated with the Internet host and a corresponding plurality of computer systems, respectively, wherein the plurality of network paths individually comprise a corresponding computer system, the Internet host, and at least one intermediate network node, and wherein the route information comprises a plurality of router labels associated with the host IP address and an intermediate network node;

extracting a location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

consulting a data store comprising at least one data set having location codes and corresponding location information;

obtaining location information from the data store corresponding to the location code associated with the one of the Internet host and the intermediate network node proximate the Internet host; and

providing a location estimate of the location of the Internet host according to the location information.

12. A method of determining the location of an Internet host using a first computer system, comprising:

measuring a first delay time relating to a first network path between a host IP address associated with the Internet host and the first computer system;

measuring a second delay time relating to a second network path between the host IP address and a second computer system;

measuring a third delay time relating to a third network path between the host IP address and a third computer system;

correlating the first, second, and third delay times; and

providing a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

13. The method of claim 12, wherein correlating the first, second, and third delay times comprises triangulating the first, second, and third delay measurements.

14. The method of claim 12, wherein correlating the first, second, and third delay times comprises:

consulting a data store comprising N sets of first, second, and third delay measurements between the first, second, and third computer systems, respectively, and N known hosts, as well as location information associated with the N known hosts, wherein N is an integer;

performing a comparison of the first, second, and third delay times with the N sets of first, second, and third delay measurements in the data store;

determining a nearest set of first, second, and third delay measurements according to the comparison; and

providing a location estimate of the Internet host according to the nearest set of first, second, and third delay measurements.

15. The method of claim 14, wherein performing the comparison of the first, second, and third delay times with the N sets of first, second, and third delay measurements in the data store comprises determining N Euclidian distances corresponding to the Euclidian distances between the N sets of first, second, and third delay measurements in the data store and the first, second, and third delay times, and wherein providing a location estimate of the Internet host according to the nearest set of first, second, and third delay measurements comprises selecting location information associated with the set of first, second, and third delay measurements in the data store associated with the smallest Euclidian distance as the location estimate.

16. The method of claim 12, wherein correlating the first, second, and third delay times comprises:

computing a first probability density function establishing a relationship between a first network delay associated with the first computer system and a first distance from the first computer system;

determining a first distance estimate representative of the distance between the first computer system and the location of the Internet host using the first delay time and the first probability density function;

computing a second probability density function establishing a relationship between a second network delay associated with the second computer system and a second distance from the second computer system;

determining a second distance estimate representative of the distance between the second computer system and the location of the Internet host using the second delay time and the second probability density function;

computing a third probability density function establishing a relationship between a third network delay associated with the third computer system and a third distance from the third computer system; and

determining a third distance estimate representative of the distance between the third computer system and the location of the Internet host using the third delay time and the third probability density function;

and wherein providing the location estimate comprises triangulating the first, second, and third distance estimates.

17. The method of claim 16, wherein determining the first, second, and third distance estimates further comprises computing an error function over a location space, and determining coordinates within the location space where the error function is minimized.

18. The method of claim 17, wherein determining coordinates within the location space where the error function is minimized comprises minimizing the error function across a list of known city locations, and wherein providing the location estimate

comprises providing the known city location corresponding with the minimum value of the error function.

19. The method of claim 18, wherein computing the error function comprises using a weighted least mean squares algorithm to optimize the location estimate.

20. The method of claim 18, wherein computing the error function comprises using a probability density estimation to optimize the location estimate.

21. The method of claim 18, wherein computing the error function comprises using a weighted least mean squares algorithm to prune a solution space, and using a probability density estimation to optimize the location estimate from the pruned solution space.

22. The method of claim 18, wherein computing the error function comprises using a probability density estimation to prune a solution space, and using a weighted least mean squares algorithm to optimize the location estimate from the pruned solution space.

23. A software tool for determining the location of an Internet host using a first computer system, comprising:

a first delay component adapted to measure a first delay time relating to a first network path between a host IP address associated with the Internet host and the first computer system;

a second delay component adapted to measure a second delay time relating to a second network path between the host IP address and a second computer system;

a third delay component adapted to measure a third delay time relating to a third network path between the host IP address and a third computer system;

a correlation component adapted to correlate the first, second, and third delay times; and

an estimation component adapted to provide a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

24. A computer-readable medium having computer-executable instructions for:

measuring a first delay time relating to a first network path between a host IP address associated with the Internet host and the first computer system;

measuring a second delay time relating to a second network path between the host IP address and a second computer system;

measuring a third delay time relating to a third network path between the host IP address and a third computer system;

correlating the first, second, and third delay times; and

providing a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

25. A system for determining the location of an Internet host, comprising:

a first component operating in a first computer system to measure a first delay time relating to a first network path between a host IP address associated with the Internet host and the first computer system;

a second component operating to measure a second delay time relating to a second network path between the host IP address and a second computer system;

a third component operating to measure a third delay time relating to a third network path between the host IP address and a third computer system;

a fourth component operating to correlate the first, second, and third delay times; and

a fifth component operating in a first computer system to provide a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

26. Geographical location estimate data associated with an Internet host, the estimate data resulting from a process of:

measuring a first delay time relating to a first network path between a host IP address associated with the Internet host and a first computer system;

measuring a second delay time relating to a second network path between the host IP address and a second computer system;

measuring a third delay time relating to a third network path between the host IP address and a third computer system;

correlating the first, second, and third delay times; and

providing a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

27. A method of determining the location of an Internet host using a first computer system, comprising:

obtaining partial IP-to-location mapping information from a data source;

obtaining network routing information;

clustering together IP addresses corresponding to hosts in the same geographic location according to network routing information to obtain cluster information;

correlating the partial IP-to-location information with the cluster information; and

providing a location estimate of the location of the Internet host according to the correlation of the partial IP-to-location information and the cluster information.

28. The method of claim 27, wherein obtaining network routing information comprises using a routing protocol.

29. The method of claim 28, wherein the routing protocol is one of BGP, RIP, OSPF, IGRP, and EGP.

30. The method of claim 27, wherein clustering together IP addresses corresponding to hosts in the same geographic location according to network routing

information to obtain cluster information comprises associating an address prefix used by a routing protocol with a geographical location.

31. The method of claim 30, further comprising:

sub-dividing the geographical location associated with the address prefix into at least two clusters according to a geographical spread associated with the geographical location.

32. The method of claim 27, further comprising sub-dividing the cluster information according to a geographical spread associated with the geographical location.

33. The method of claim 27, further comprising computing a dispersion metric representative of the accuracy of the location estimate of the location of the Internet host.

34. The method of claim 33, further comprising selectively providing the location estimate if the dispersion metric is less than a threshold value.

35. The method of claim 27, further comprising:

obtaining route information relating to a first network path between a host IP address associated with the Internet host and the first computer system, wherein the first network path comprises the first computer system, the Internet host, and at least one intermediate network node, and wherein the route information comprises a plurality of router labels associated with the host IP address and the at least one intermediate network node;

extracting a first location code from the route information corresponding to a router label associated with one of the Internet host and an intermediate network node proximate the Internet host;

consulting a data store comprising at least one data set having location codes and corresponding location information;

obtaining first location information from the data store corresponding to the first location code associated with the one of the Internet host and the intermediate network node proximate the Internet host; and

providing a first location estimate of the location of the Internet host according to the first location information from the data store corresponding to the first location code.

36. The method of claim 35, further comprising:

measuring a first delay time relating to a first network path between a host IP address associated with the Internet host and the first computer system;

measuring a second delay time relating to a second network path between the host IP address and a second computer system;

measuring a third delay time relating to a third network path between the host IP address and a third computer system;

correlating the first, second, and third delay times; and

providing a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

37. The method of claim 27, further comprising:

measuring a first delay time relating to a first network path between a host IP address associated with the Internet host and the first computer system;

measuring a second delay time relating to a second network path between the host IP address and a second computer system;

measuring a third delay time relating to a third network path between the host IP address and a third computer system;

correlating the first, second, and third delay times; and

providing a location estimate of the location of the Internet host according to the correlation of the first, second, and third delay times.

38. A software tool for determining the location of an Internet host using a first computer system, comprising:

a first component adapted to obtain partial IP-to-location mapping information from a data source;

a routing protocol component adapted to obtain network routing information;

a clustering component adapted to cluster together IP addresses corresponding to hosts in the same geographic location according to network routing information to obtain cluster information;

a correlation component adapted to correlate the partial IP-to-location information with the cluster information; and

an estimation component adapted to provide a location estimate of the location of the Internet host according to the correlation of the partial IP-to-location information and the cluster information.

39. A computer-readable medium having computer-executable instructions for:

obtaining partial IP-to-location mapping information from a data source;

obtaining network routing information;

clustering together IP addresses corresponding to hosts in the same geographic location according to network routing information to obtain cluster information;

correlating the partial IP-to-location information with the cluster information; and

providing a location estimate of the location of the Internet host according to the correlation of the partial IP-to-location information and the cluster information.

40. A system for determining the location of an Internet host, comprising:

a first component operating to obtain partial IP-to-location mapping information from a data source;

a second component operating to obtain network routing information;

a third component operating to cluster together IP addresses corresponding to hosts in the same geographic location according to network routing information to obtain cluster information;

a fourth component operating to correlate the partial IP-to-location information with the cluster information; and

a fifth component operating to provide a location estimate of the location of the Internet host according to the correlation of the partial IP-to-location information and the cluster information.

41. Geographical location estimate data associated with an Internet host, the estimate data resulting from a process of:

- obtaining partial IP-to-location mapping information from a data source;
- obtaining network routing information;
- clustering together IP addresses corresponding to hosts in the same geographic location according to network routing information to obtain cluster information;
- correlating the partial IP-to-location information with the cluster information; and
- providing a location estimate of the location of the Internet host according to the correlation of the partial IP-to-location information and the cluster information.